(pre-print version of a manuscript published in Elsevier Microprocessors and Microsystems)

Protection of heterogeneous architectures on FPGAs: An approach based on hardware firewalls

Pascal Cotret [α], Guy Gogniat [β], Martha Johanna Sepúlveda Flórez [β,γ]

[α] IETR/SCEE, CentraleSupélec - FRANCE

pascal.cotret@centralesupelec.fr

[β] Lab-STICC Laboratory, University of South Brittany - FRANCE

guy.gogniat@univ-ubs.fr

[β,γ] Institute for Security in Information Technology, Technical University of Munich - GERMANY

johanna.sepulveda@tum.de

arXiv:1602.05106v1 [cs.CR] 16 Feb 2016

Abstract

Embedded systems are parts of our daily life and used in many fields. They can be found in smartphones or in modern cars including GPS, light/rain sensors and other electronic assistance mechanisms. These systems may handle sensitive data (such as credit card numbers, critical information about the host system and so on) which must be protected against external attacks as these data may be transmitted through a communication link where attackers can connect to extract sensitive information or inject malicious code within the system. This work presents an approach to protect communications in multiprocessor architectures. This approach is based on hardware security enhancements acting as firewalls. These firewalls filter all data going through the system communication bus and an additional flexible cryptographic block aims to protect external memory from attacks. Benefits of our approach are demonstrated using a case study and some custom software applications implemented in a Field-Programmable Gate Array (FPGA). Firewalls implemented in the target architecture allow getting a low-latency security layer with flexible cryptographic features. To illustrate the benefit of such a solution, implementations are discussed for different MPSoCs implemented on Xilinx Virtex-6 FPGAs. Results demonstrate a reduction up to 33% in terms of latency overhead compared to existing efforts.

1 Introduction

For many years, embedded systems are used in our daily life: we found them in electronic devices, automotive applications, telecommunications systems and so on. When designing such systems, several issues have to be taken into account and one of the major concerns is about security. Since the late 90s, security has become a key point in the development of embedded systems [1]. The number of weaknesses is in constant progress and electronic devices have to process data with various security requirements. According to [2], security criteria are communications security, storage security, inputs/outputs security and users authentication. This work focuses on the two first criteria (communications and storage).

First of all, this work considers communication protection as a key point in embedded systems development as communications channels convey several data types (application codes, confidential data, cryptographic elements and so on) with various needs in terms of security: confidential data must not be revealed to an unauthenticated user while application may be accessible through a specific interface (for instance, for development purposes). Then, this work also takes care of data

storage security: memory elements are another critical entry point for attackers as they potentially contain plaintext data.

This work is organized as follows. Section 2 presents related works and our constraints regarding the architecture. Sections 3 and 4 describe our solution in a static and dynamic approach. Section 5 gives an analysis in terms of security and provides implementation results in comparison with other approaches.

2 Scientific context

2.1 Related works

Several studies dealing with security in embedded systems have been published [1, 2]. Security mechanisms can be implemented in two ways: hardware blocks or software functions. Software solutions are generally slower, in terms of latency, than a pure hardware-implemented security solution. Furthermore software solutions are generally more easily compromised than hardware countermeasures. In this section, several works about memory protection are presented. Then, regarding internal transactions protection, an overview of the main solutions is proposed.

2.1.1 Memory protection approaches

In order to provide countermeasures against the threat model defined in Section 2.2, a key point is to address memory protection. An obvious solution is to implement cryptographic features for memory confidentiality and integrity. XOM [3] is a solution mixing confidentiality and integrity for systems where the external memory can be tampered. The implementation requires adding hardware modules and modifying the processor structure. Using such a solution, performances are quite spoiled as authors [3] announce a 50% loss. AEGIS architecture [4] is another approach based on a security-enhanced processor embedding confidentiality and integrity functions. Depending on processor configuration (cache size), memory slowdown is between 3.8% and 130%. Bossuet et al. [5] made an in-depth comparison of existing cryptographic processors where some of them, such as HCrypt, were implemented on FPGAs. Some of these processors are efficient but do not cover our threat model defined in a further section. In [6] authors describe a solution (called SecSoft) to protect software with a hardware Encryption Management Unit (EMU). This work proposes a latency analysis of several modes (block/counter modes for encryption function and with/without encryption). Latency overhead on a ML301 platform goes from less than 10% (block mode, unencrypted) up to 80% (block mode, encrypted). This solution does not provide mechanisms targeting integrity. [7] proposes the PE-ICE solution to check integrity in parallel to encryption (i.e. confidentiality). The worst case implementation shows a performance loss of 20% and for a confidentiality only implementation, a 4% loss is given. Vaslin et al. [8] proposes a confidentiality and integrity hardware block based on AES for confidentiality and cyclic redundancy check for integrity, performance loss is about 13-14%. In [9], authors extended this work by using the AES-GCM algorithm (this option is also used in this work), integrity is done by a low latency function [10]. Other methods such as hash trees and formal verification [11, 12] are use to protect memory contents.

Another solution is to use the built-in MMU (Memory Management Unit) available in some processors. This work is implemented on a Xilinx FPGA where the softcore Microblaze is provided as a general purpose processor. Microblaze MMU [13] provides a simple access control allowing to get read-only or full-access memory pages in the system, this control can be disabled in a configuration register.

All these works propose solutions to provide encryption methods for external memories protection with different performance versus security tradeoff. In order to protect the target system from attacks on the external memory, a trivial solution consists in building a fully-protected external

memory unit. Unfortunately, in this case, each memory access implies a ciphering/deciphering latency penalty. Thus, such an approach is strongly penalizing regarding the overall latency overhead of an application. To mitigate this point, our work proposes to implement cryptographic features only on specific memory pages, defined by application requirements, avoiding such a systematic latency penalty. Therefore, some pages are still not protected, that is the reason why internal traffic protection and/or monitoring must be also addressed in the context of embedded systems security.

2.1.2 Bus and NoC-based security methods

Regarding large scale systems with NoC-based communication architecture, [14] proposes a solution where security controls are done in each network interface in a distributed manner. In this case, a security manager unit gathers individual interfaces information and performs countermeasures and security updates (done through dynamic partial reconfiguration). This method takes into account processor facing denial of service attacks but does not offer ciphering function (however, authentication and integrity features are available). Fiorin et al. [15, 16, 17] propose an alternative to this approach providing ?security sensors? inside network interfaces to refine controls (NI in Figure 1). These sensors are able to block incoming malicious data when parameters are not proven. These parameters are stored in a trusted CAM (Context-Addressable Memory). Finally, a SNM (Security Network Manager) gathers information from NIs to detect potential collisions and errors in data traffic. Unlike [14], security mechanisms can be updated without partial reconfiguration of the FPGA chip, update is done using memory rewriting. [14] and [15, 16, 17] do not offer cryptographic features to cipher data transmitted in the communication network.

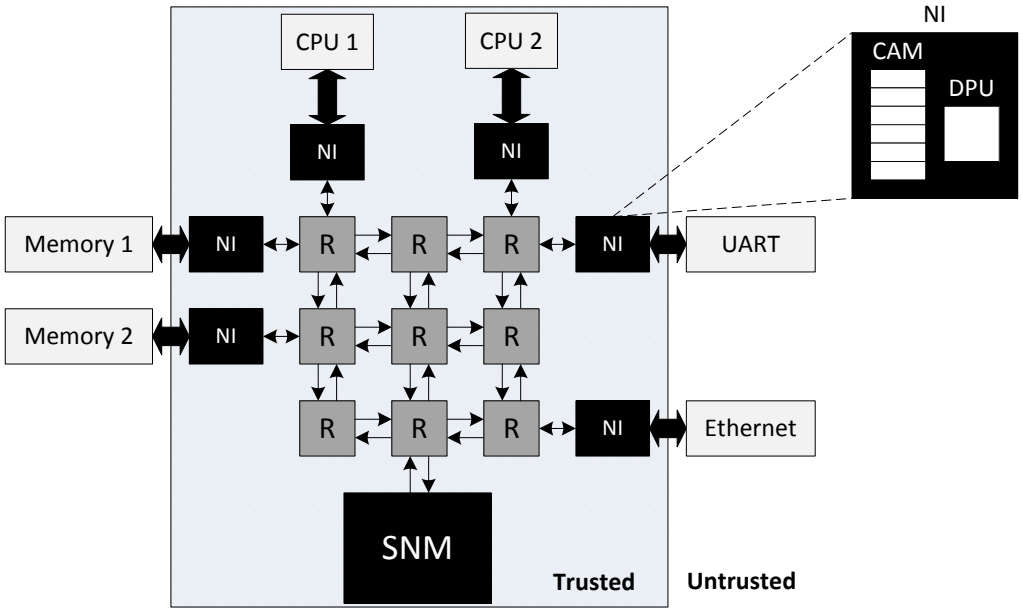

Figure 1: Fiorin's approach [17, 15]

The solution proposed by Fiorin et al. is based on a secured Network Interface (NI) with a DPU mechanism (Data Protection Unit, see Figure 1). This mechanism allows or not a transaction according to parameters stored in individual trusted CAM memories. [15] requirements aim to cover a threat model with denial of service attacks but cannot protect systems against IP modifications performed by an attacker. This distributed approach has a low-latency and also presents some robustness. Controls are done in each interface, even if one of them is corrupted, other interfaces should still continue to work. Regarding NoC-based security solutions, we can also cite [18, 19]. Structure proposed by Sepulveda et al. is based on a hierarchical NoC with low NoCs (each low NoC is a sub-network having a single security policy) and a single high NoC acting as a global security

manager (connections with each low NoC). Regarding NoC-based security, LeMay et al. [20] propose a mechanism to detect abnormal behaviors in a NoC protocol when a malicious IP is inserted in the system architecture: the solution based on AXI signals has an area overhead up to 23%.

Then, for bus-based MPSoCs, the main contribution was published by Coburn et al. [21]. This approach is similar to [15, 16, 17] as it is based on security-enhanced network interfaces called SEI (*Security Enforcement Interface*) but in a centralized approach. The main drawback of this solution is that security information is sent to a global security manager which is the only component able to perform controls. Therefore, latency overhead is increased (see Section 5). This solution does not offer security updates or cryptographic features. Coburn et al. approach suggests to centralize all the controls in a single module; therefore, as soon as this module is corrupted, system security is compromised.

Compared to these efforts, our approach provides a distributed solution with update mechanisms. We are able to dynamically adapt the security policies based on the instantaneous threats. We also provide some cryptography mechanisms in order to fully protect the system. We propose a low-latency solution in order to reduce the performance penalty due to security.

2.1.3 Other methods

Other solutions propose a physical separation of components in order to define secured and non-secured areas. For instance, [22] use *moats* during place-and-route, routing is forbidden in these areas, it allows isolating specific areas of the FPGA. Other contributions propose to combine hardware and software elements to provide security. [23, 24] focus on virtualization, an hardware mechanism associated with software services that manages several software tasks in a secured manner on an heterogeneous architecture. Another solution is proposed by ARM, Trustzone [25]. It offers the possibility to have two areas with different privileges. A hardware module is in charge of monitoring data, communications between these two areas are secured if the parameters of the architecture are set to secure. However, virtualization-based solutions are not adapted to our context: for small-scale MPSoCs, it is assumed that hardware-based solutions are desirable. Furthermore, this work aims to provide a solution with the lowest impact on existing OS or architecture: as virtualization adds some mechanisms to both software and hardware layers, these solutions are not considered in our further comparison. Our contribution could be extended with mechanisms as provided by [22]. Our approach goes into the same direction as works developed by [23, 24] and [25], but we propose a more comprehensive solution providing cryptography, filtering and monitoring.

2.2 Platform characterization and threat model

This work implies a set of constraints about architectures where our approach can be applied:

- FPGA chips are chosen as they allow short development times and system reconfigurability whereas ASICs cannot be reprogrammed.
- This work is done in a context of a project where applications do not require a large number of IPs (small-scale up to medium-scale architectures). Therefore, it is assumed that a single communication bus is able to manage all case studies targeting MPSoCs.
- Regarding communication protocols, AXI protocol from ARM [26] is chosen. In fact, this protocol is the standard in latest Xilinx development tools and should provide a compatibility with other ARM-based technologies in the future such as Zynq circuits (embedding an FPGA and a Cortex-A9 processor [27]) or Armadeus [28].

This work is based on a specific threat model for FPGA-based MPSoC architectures. At a high level of abstraction, it is assumed that the FPGA itself is a trusted component (i.e. attackers cannot

directly access it). On the opposite, both external bus and external memory are considered as untrusted areas. As a consequence, we can define a range of attack scenarios that must be covered by our approach. On the one hand, following scenarios are taken into account (non exhaustive list):

- Attack A1. External memory is compromised by replacement or modification. In this scenario, it is considered that the attacker is able to physically replace the external memory by its own chip. The attacker is also able to modify its contents by any means: accesses can be logical (access with another processor) or even physical (for instance, if memory contents are modified by overheating or exposure to an electromagnetic field).

- Attack A2. Snooping on-chip to off-chip bus through an external probe. The attacker is able to access the physical bus between the FPGA chip and the external memory. The attacker can read data on the bus or inject malicious data.

- Attack A3. Internal snooping on on-chip busses for data values by an untrusted IP included in the SoC. It is considered that the attacker is able to implement an IP (hardcore or softcore) aiming to add data to the on-chip traffic of the system. These data can be both packets with valid address or denial-of-service attempts. The attacker is also able to retrieve data from the bus for example to get private information.

On the other hand, following scenarios are not considered in our threat model:

- Snooping on the on-chip busses through package modification or physical probe. This work considers that the chip cannot be physically modified after bitstream downloading.

- Attempting to measure the contention on the bus or cache-timing attacks.

- Other side-channel attacks: power analysis, electromagnetic analysis and so on.

These scenarios can be combined to construct complex attack scenarios. Two scenarios will be tested in Section 5.2 to compare this work with existing techniques. As countermeasures, this work proposes to focus on protecting and monitoring the internal traffic as well as protecting the external memory.

2.3 Main contributions

Contributions provided in this work aim to protect external memories and internal traffic in multi-processor architectures implemented on FPGA circuits. Our approach provides the demonstration that a complete end-to-end solution is possible. Such an approach allows reducing the performance penalty due to cryptographic features as all communications are not encrypted while still guaranteeing that unprotected code or data will not lead to an attack on the system thanks to internal traffic protection/monitoring. Our solution represents an interesting alternative compared to [15, 21] with an additional security layer including cryptographic services.
Table 1 presents a summary of qualitative parameters targeted by our solution (which is basically described in Section 3). This work enhances existing efforts by providing a solution for a comprehensive protection of internal and external communications:

- For traffic monitoring and protection, our approach is similar to [15] but it provides a more efficient solution for security updates in terms of memory consumption. This contribution is detailed in Section 5.

- For memory countermeasures, our approach is similar to [6, 7] but the AES-GCM cryptographic primitive is used as proposed by [9]. Such a solution allows reducing the cryptography latency overhead while maintaining a strong level of protection. Compared to [6, 7], we rely on a more comprehensive security policy where developer can finely tune his/her cryptography needs.

Table 1: Comparison of existing solutions

		[9]	[13]	[3, 4]	[14]	[16]	[19]	[21]	This work
Communication technology		N/A	N/A	N/A	NoC (mesh)	NoC (mesh)	NoC (hierarchical)	Bus (AMBA)	Bus (AXI)
Approach		Mem. extension	Proc. extension	Proc. extension	Centralized IP	Distributed interfaces	Distributed IPs	Centralized IP	Distributed interfaces
Security updates		No	No	No	No	Yes	No	No	Yes
Counter-measures	**Confidentiality**	Yes	No	Yes	No	No	No	No	Yes
	Integrity	Yes	No	Yes	No	No	No	No	Yes
	Memory partitioning	Yes	Yes	No	No	Yes	Yes	Yes	Yes
	Traffic monitoring	No	No	No	Yes	Yes	Yes	Yes	Yes
	Access control (R/W)	No	Yes	No	Yes	Yes	Yes	Yes	Yes
Threat model coverage	**Attack A1**	Yes	No	Yes	No	No	No	No	Yes
	Attack A2	Yes	No	Yes	No	No	No	No	Yes
	Attack A3	No	No	No	Yes	Yes	Yes	Yes	Yes

Then, this work also presents a feedback feature for security mechanisms (Section 4): that is to say how attacks are detected and how the update manager takes care of these information. Finally, an evaluation in terms of implementation results and security analysis is given.

3 Static security for communications and memories

In order to secure MPSoCs, this work proposes an approach with hardware security components (also known as "firewalls") implemented in the architecture. Figure 2 shows an example of such a system.

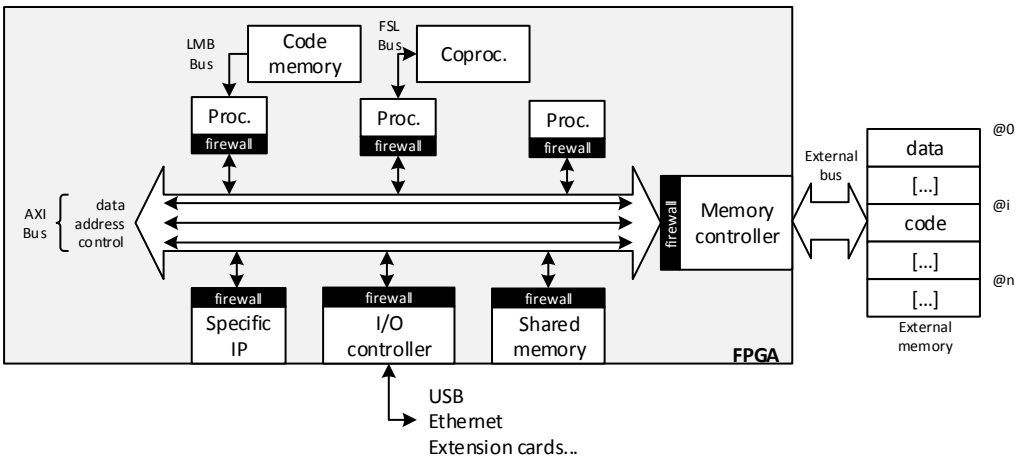

Figure 2: Generic MPSoC with hardware firewalls

A key point when dealing with security corresponds to the need of clearly setting up all parameters chosen to protect an MPSoC against a given threat model. Thus before discussing our approach, a set of rules (or security policies) is defined for each access in the system.

3.1 Security policy

Each security policy is defined by several parameters:

- The memory segment base/high addresses. An essential feature of firewalls is to block illegal accesses in terms of addresses. For instance, an illegal access to a memory section by a general purpose processor which is not allowed by the application specifications must be discarded (writing a confidential data 190 into memory, reading a section with cryptographic-related information...).

- In order to avoid attacks such as "buffer overflow", data format has to be analyzed for each transaction (both read and write operations).
- As systems contain an external memory, cryptographic primitives are used for data protection. According to the application requirements, some memory sections may not be ciphered if the contents are not critical.

As described in Figure 2, a generic MPSoC contains an external memory. Thus the most radical solution to guarantee security consists in protecting the whole memory in terms of confidentiality and authentication. In such a case, an attacker cannot decipher or modify its contents. Unfortunately, this solution has an important overhead in terms of latency (see Section 5). An alternative solution consists in ciphering only the most critical memory sections. In this case, an attacker can read and write all the other plaintext memory sections. That is the reason why implementation of a cryptographic function only is not efficient enough to protect the target MPSoC against the defined threat model. Therefore, firewalls also aim to protect the system against these plaintext contents by monitoring the internal communications.

3.2 Local Firewall

Each firewall is composed of several blocks (*Security Builder* and *Firewall Interface*) connected as defined in Figure 3. The *Security Builder* module is responsible for security policy management and the *Firewall Interface* is a mandatory checkpoint towards the external world (that is to say, the system communication bus). This module acts as a gateway between the AXI-4 communication bus and the associated IP (it can be dedicated to an application, an I/O controller, a memory interface...).

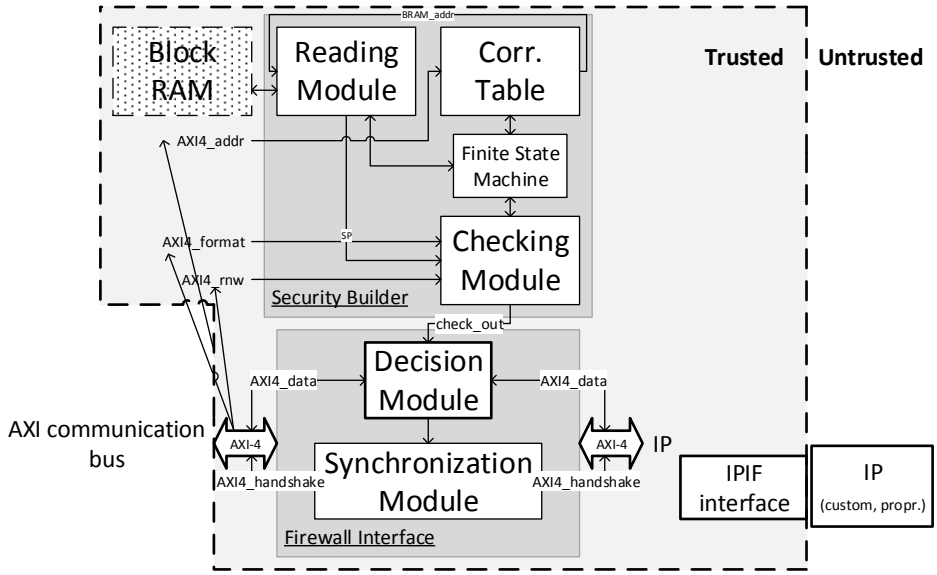

Figure 3: Structure of a Local Firewall

The whole firewall structure is considered as a trusted area (it is implemented in hardware in the FPGA chip), the only part to be untrusted is composed of the external bus and the external memory (components linked to the *Cryptographic Firewall*).

3.2.1 Firewall Interface

The *Firewall Interface* module (#1 in Figure 3) performs mainly two tasks:

- Once a data block has been analyzed and validated (its parameters match with the associated security policy), the *Firewall Interface* transmits data to the transaction target (communication bus or IP); this operation is done by the *Decision Module*.
- The Firewall Interface synchronizes the communication protocol signals such as handshake signals (*AXI_WVALID*, *AXI_RREADY*...) or other control signals (*AXI_WSTRB*, *AXI_WLAST*...) to keep valid transactions. Data block is synchronized with its control signals in order to avoid data loss and/or duplication. The sub-component *Synchronization Module* is composed of a set of flip-flops (see Figure 4) where the clock port (at least its rising edge) is the acknowledgment signal *check_out* sent by the *Security Builder* when all the data checking has been performed.

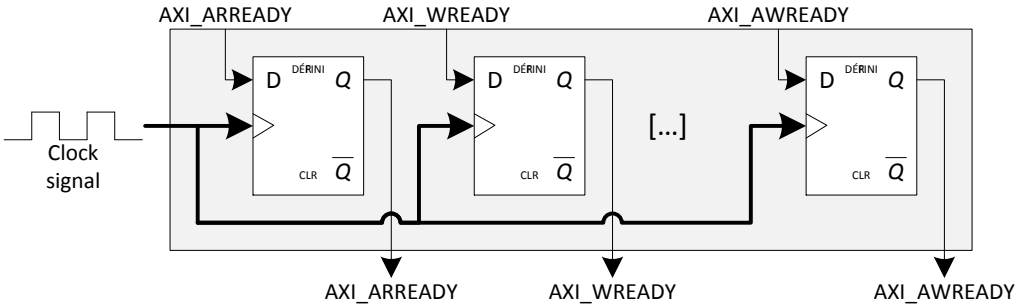

Figure 4: Structure of a Synchronization Module

As all the flip-flops are connected in parallel, the overall *Firewall Interface* latency is of 2 cycles for each data block: 1 cycle for the *Decision Module* and 1 cycle for the synchronization step.

3.2.2 Security Builder

The *Security Builder* module is the main component within firewalls (#2 in Figure 3). It is composed of 4 submodules:

- **Correspondence Table (CorrTable)**. Security policies are stored in a Block RAM memory considered as a trusted entity (upper left corner in Figure 3) and can be identified by an address. Internal structure of *Correspondence Table* is detailed in Figure 5. Each policy is defined for a given physical address space (with lower and upper bounds), transcription towards policies addresses is done as defined in Algorithm 1.

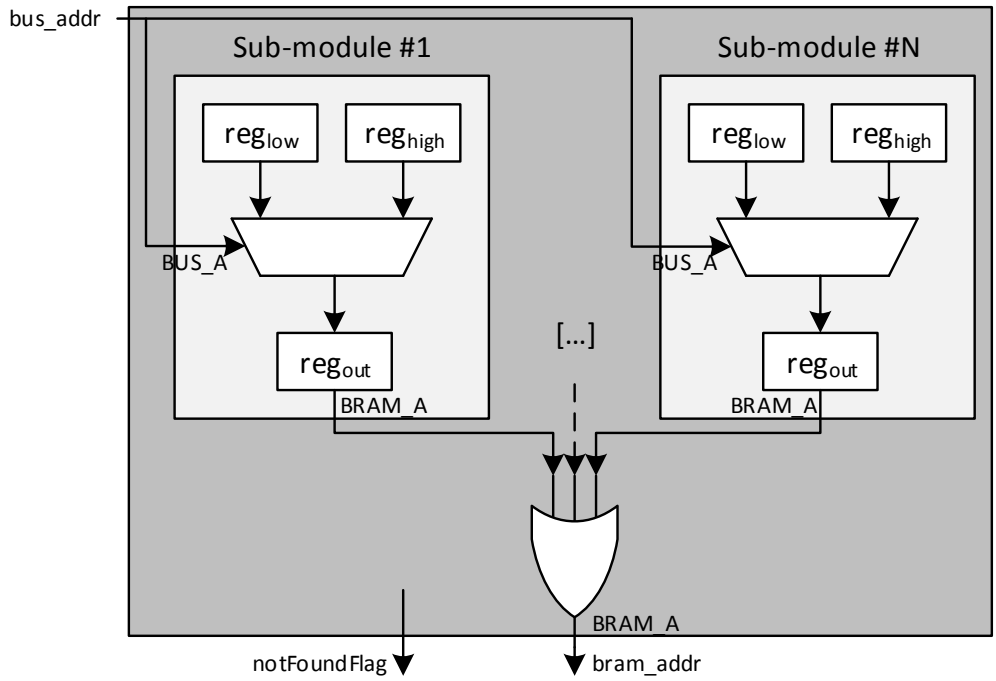

Figure 5: Implementation of a Correspondence Table

Given that N is the number of entries of the *Correspondence Table* (i.e. the number of security policies embedded in the firewall), $BRAM_A$ is initially equal to zero. If the bus address is contained in a given address space (defined with reg_{low} and reg_{high}), $BRAM_A$ is set to this value. Otherwise if BUS_A is outside the address space, a flag is set to high value (this flag reports an error and is treated by the global finite state machine embedded in each firewall).

Algorithm 1: Correspondence Table loookup

Require: $n \leq N_{max}$ {N_{max}: `max number of entries`}
1: $BUS_A \leftarrow bus_addr$ {`Get a copy of bus address`}
2: $BRAM_A \leftarrow 0x00000000$ {`Initialize output`}
3: **for** $SUBMODULE = 1$ to N **do**{`For each submodule`}
4: **if** $BUS_A \in [reg_{low}; reg_{high}[$ **then** {`If submodule contains the bus address`}
5: $BRAM_A \leftarrow reg_{out}$ {`Write the location of security policy`}
6: **end if**
7: **end for**
8: **if** $BRAM_A = 0$ **then** {`If the output register is null`}
9: $not_found_flag \leftarrow 1$ {`Set a warning flag !`}
10: **end if**
11: $bram_addr \leftarrow BRAM_A$ {`Export the output to other components`}

- **Reading Module (ReadMod)**. *ReadMod* is in charge of reading security policies from the dedicated Block RAM trusted memory and extracting security parameters to be transmitted to the *Checking Module*. For a single 32-bit word, a reading buffer is filled in 1 clock cycle. Then, security parameters are transmitted to the *Checking Module* in 1 additional cycle.

- **Checking Module (CheckMod)**. For a simple implementation, *Checking Module* verifies two parameters: read/write access right and data format thanks to a set of comparators and a test function (see Figure 6).

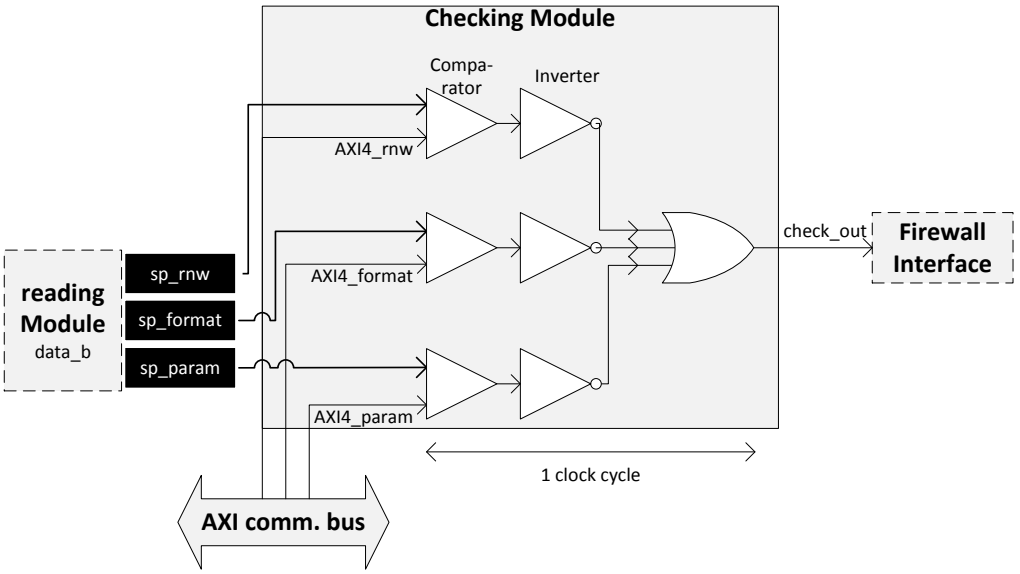

Figure 6: Structure of a Checking Module

First, there is a preliminary test on the value of the *ARID* signal extracted from the bus: when this signal is not equal to 0, the data being analyzed is associated with a read; otherwise (*ARID=0*), it is related with a write operation. This test function output is used to check two parameters:

- The select input of a multiplexer for format value verification (*ARSIZE* and *AWSIZE* contain format values for read and write operations [26]).
- Then, access right and data format are compared with *sp_rnw* and *sp_format* values extracted from the security policy. As all the parameters are instantiated in parallel (see Figure 6), all the parameters are analyzed at the same time. In Figure 6, inverters and the OR gate produce a global signal representative of all the controls done in the *Checking Module*: if one or more comparators go wrong, the final output will be 0 (in the other case, when all comparators go fine, the output is 1).

- **Finite State Machine (FSM)**: the FSM managing firewall behavior.

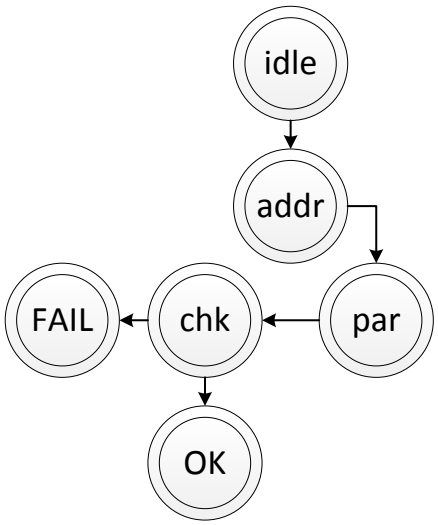

Figure 7: FSM associated with the Security Builder module

FSM described in Figure 7 represents the simplified behavior of a *Security Builder* embedded in a *Local Firewall* (see Figure 3). It contains the main states of this module (without error procedure or cryptographic feature):

- *Idle* state. The *Security Builder* module waits for a new incoming data to be analyzed (and received from the *Firewall Interface*).
- *Addr* state. Allowance to recover Block RAM address where data is located.
- *Par* state. Security policy reading step
- *Chk* state. Checking operation performed in 2 clock cycles: 1 cycle for the preliminary test and another cycle for the combinatorial part (comparators and gates). If parameters match with the security policy, *OK* state is enabled (otherwise, *FAIL* state).

In order to have a complete firewall behavior, the final step is the transmission of the *check_out* signal to the *Firewall Interface* sub-module: if checking operations match with the security policy, signals are transmitted in a synchronized manner; otherwise, data are not transmitted (the process resulting from an error is described later in this paper). Checking a 32-bit data word takes 4 clock cycles. Therefore, checking N 32-bit data words is done in 4N clock cycles in a non-pipelined architecture.

Local Firewalls allow to monitor and to filter any communication on the bus. They do not monitor the internal execution flow of processors and their associated cache behaviors. If an attacker sets up a cache based attack (e.g. trace-driven or time-driven), *Local Firewalls* will not block such an attack as it does not correspond to illegal operations (i.e. access right, data format). Additional solution (e.g. software-based solution) should be used to provide some countermeasures against this type of attacks. If an attacker tampers an application code that is not protected with integrity mechanisms and if this malicious code tries to perform illegal accesses then *Local Firewalls* will detect and block this attack. *Local Firewalls* do not protect against side channel attacks but allow to detect and stop any attack that does not respect the security policies associated with an application. Extending these mechanisms with cryptography properties enables to address a large class of potential attacks.

3.3 Cryptographic Firewall

The external memory controller has a dedicated firewall offering cryptographic features in addition to security services embedded in a *Local Firewall*. This type of firewall (also known as *Cryptographic Firewall*) is able to protect data stored in the external memory in terms of confidentiality and authentication. The overall structure of a *Cryptographic Firewall* is shown in Figure 8.
The main difference between a *Local* and a *Cryptographic Firewall* is the *Security Builder*. Datapath is modified to take into account confidentiality and integrity features implemented in a *Cryptographic Firewall*:

- In case of a read: the first step is to get the security policy associated with data currently analyzed. Then, data is processed by the cryptographic module (deciphering and eventually authentication checking) before being sent to the *Security Builder* (verification of access rights, data format...).
- In case of a write: datapath goes through the *Security Builder* before going to the cryptographic module if parameters match (*check_out=1*).

In this contribution, firewalls have to operate with a low-latency feature. Therefore, a global protection of external memories (in terms of confidentiality and authentication) is not the best option.

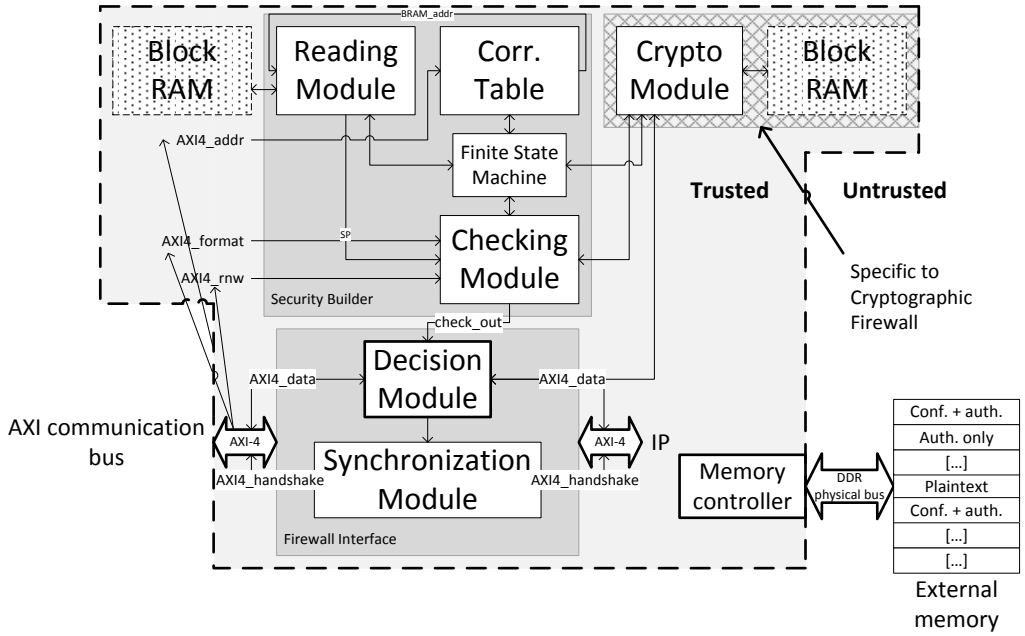

Figure 8: Structure of a Cryptographic Firewall

That is the reason why we want to implement an algorithm offering confidentiality and authentication separately (the security mode is specified by the security policies). Four options are considered in this work:

- AES + MD5. The simplest solution uses the AES standard for confidentiality and a hash function such as MD5 for integrity. Even if it has been developed in the early 90s, MD5 is still widely used [29, 30].

- AES + SHA-2. Developed when security weaknesses were discovered in MD5. According to several surveys[1], other weaknesses were highlighted and implementations with better security were developed.

- AES + SHA-3. NIST[2] launched a contest in 2007 to define SHA-3. The winner was announced in October 2012, it is a protocol developed by STMicroelectronics and NXP named Keccak[3]. For experimentations, we use results extracted from the ATHENA database created by George Mason University[4] for implementations on Xilinx Virtex-6 FPGAs.

- AES-GCM. This particular mode of AES is able to perform authenticated ciphering. Confidentiality is guaranteed by plaintext ciphering in CTR mode and authentication is performed by a MAC calculation. [9] proposes a comparison of various AES modes and a detailed analysis of AES-GCM mode.

In order to choose the best option in our context of low latency, two metrics are used: latency on a 32-bit data block and throughput of a basic implementation. Results associated to each option are shown in Table 2.
According to Table 2, the best option is AES-GCM (four times faster than solutions based on hash functions). The structure of this mode is described in Figure 9. According to these first results, Keccak (SHA-3 standard) has also interesting performances and could be considered as another alternative.

[1] http://eprint.iacr.org/2004/207.pdf
[2] National Institute of Standards and Technology
[3] http://keccak.noekeon.org/
[4] http://cryptography.gmu.edu/athenadb/fpga_hash/table_view

Table 2: Modes comparison for the cryptographic block

	Latency (# of cycles)	**Throughput**
AES + MD5 [29, 30]	90	up to 725 Mbits/s
AES + SHA-2 [31]	74	up to 1.8 Gbits/s
AES + SHA-3 (Keccak)	25	around 30 Gbits/s
AES-GCM [32]	25	30 Gbits/s

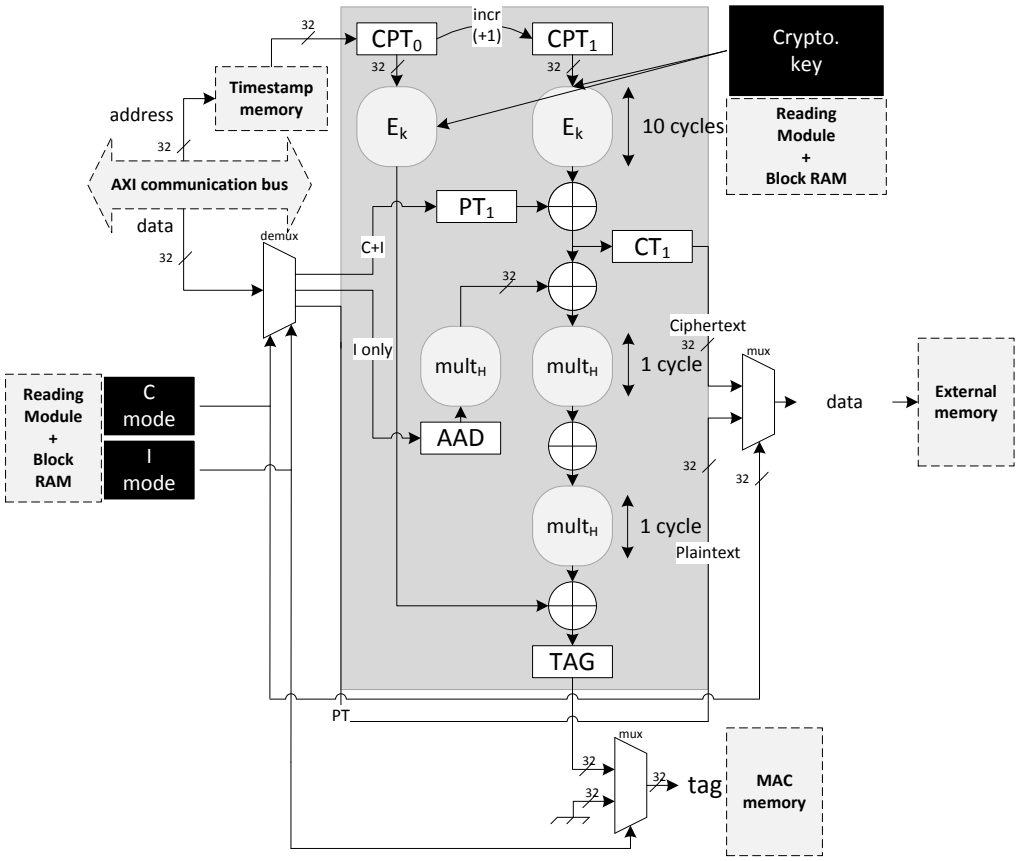

Figure 9: Cryptographic Module: AES-GCM

According to Figure 9, a single module is able to perform several options ("confidentiality and authentication", "authentication only"or "plaintext"). Multiplexors and demultiplexors ports are implemented in order to take into account the security policy parameters defining the confidentiality and authentication properties (respectively *Cmode* and *Imode*). It allows the block to take different datapaths producing ciphertext and tag (and even bypass the whole module if no cryptographic feature is enabled). Keys used for encryption/decryption are extracted from the security policies and sent by the *Reading Module* (as all information related to MACs).

In AES-GCM algorithm, confidentiality is performed by plaintext ciphering. AES function (in CTR mode) in *Ek* blocks generates keystreams (with the key extracted from the security policies stored in a trusted on-chip memory). A timestamp value is used as an input of the AES-GCM to avoid replay attacks. Finally, a XOR is performed with the keystream and the plaintext to produce the ciphertext which can be transmitted to the external memory controller. Authentication is made

with a MAC based on universal hash. AAD is a data that can potentially be authenticated without being ciphered.

In Figure 9, encryption of a 32-bit data (*Ek* uses a 128-bit key and a 128-bit data containing the data block to be ciphered on 32 bits and padded with zeros) is performed in 10 clock cycles and authentication in 2 additional cycles (2 Galois multipliers MULTH are needed [9, 10]). The overall latency for a set of N 32-bit data protected in terms of confidentiality and authentication is given by Equation 1.

$$latency(N) = 10 + (10 + 2) * N \quad (1)$$

Finally, using AES-GCM allows firewalls to provide low-latency cryptographic features to the final user while keeping flexibility in terms of available modes (confidentiality and authentication). In particular for authentication, AES-GCM takes 2 clock cycles while MD5 (respectively SHA-2) takes 64 (respectively 80) cycles to do so. Tag produced by the AES-GCM core is not ciphered as it is stored in a trusted Block RAM memory. As Block RAMs are dual-port memories, one port is left free for further parameters update operations (these operations are described in Section 4).

4 Update services for dynamic hardware firewall configuration

4.1 Security leakages in static security

Section 3 proposes a solution to protect a multiprocessor architecture with static security enhancements. Unfortunately, this is a solution where security policies cannot be updated in case of an attack. At this step, firewalls only detect an attack thanks to the *Security Builder*. Therefore, mechanisms are needed to update firewall rules without creating security weaknesses within the system. Security update can be done by partial or complete reconfiguration (download of a partial or complete bitstream). However, the solution proposed to update the security policies in this section aims to have the following features:

- No system blocking. System behavior should not be stopped during security updates. Furthermore, malicious data must not leak during this process.
- Low-latency update. Security must be updated as fast as possible.
- Different security modes. Firewalls should offer a hierarchy of security levels in order to allow mechanisms to be set in a more or less permissive mode. It is assumed that security levels are defined by access rights (read/write, read-only and so on).

4.2 Security policies evolution

Two categories of components are defined according to their ability to handle confidential information. For instance, critical IPs such as encryption blocks must not reveal information in case of an attack. If an attacker is able to extract cryptographic keys, it would be a major failure for the embedded system. In this case, critical IPs must be placed in a quarantine mode where no read or write are allowed. For non-critical IPs, an intermediate mode is defined where only read accesses are allowed. This feature allows backing up firewalls configuration and confidential information before eventually switching to the quarantine mode.

In the most critical case, when an attack is detected by a firewall already set to the quarantine mode, it is assumed that the system must be completely reset. Therefore, the initial bitstream (containing initial security policies) is downloaded into the FPGA chip to ensure a secure execution environment to the system.

When a firewall has to be updated, there are two potential areas to be modified:

- *Correspondence Table* module: it defines the relationship between the address spaces and a related security policy. In this work, it is considered that all the address spaces of each IP are covered by, at least, one security policy. Therefore, the only component to be updated is the memory containing the security policies as explained below.
- Block RAM memory containing the security policies. These memories contain the rules to be verified for each communication and memory access.

For both *Local* and *Cryptographic* firewalls, information related to read/write access rights is stored in a 32-bit word in a Block RAM memory. As a result, updating a security policy means writing a 32-bit word (with new access rights values) in a BRAM; this is done in one clock cycle[5]. Finally, updating N security policies in one firewall is performed according to the following equation:

$$duration(N) = N\ cycles \tag{2}$$

4.3 Update services architecture

When an attack is detected, Block RAMs must be updated with the new security policies in order to keep a secure execution environment. All the components are connected to an AXI-Lite communication bus (labeled "security bus" in Figure 10). A dedicated processor (labeled "update processor" in Figure 10) keeps records of all important events in a log file (timestamps, attacks and update progress).

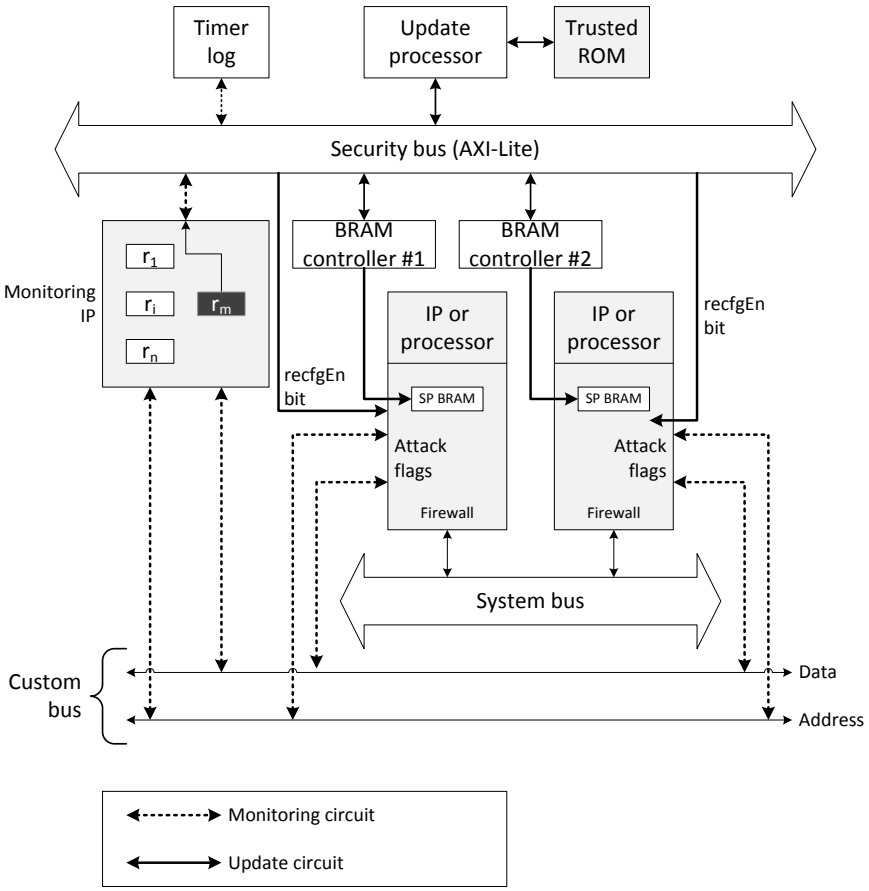

Figure 10: Global architecture of the monitoring area and update services

[5]http://www.xilinx.com/support/documentation/ip_documentation/axi_bram_ctrl/v1_03_a/ds777_axi_bram_ctrl.pdf

Each firewall has a connection with a monitoring IP through a dedicated bus. While this IP and a timer are used for attack detection and reporting, BRAM controllers are used to update the security policies embedded in each firewall BRAM memory. As the code of the "update processor"is stored in a trusted on-chip memory (as the architecture is composed of the security bus, timer log and custom bus), this update process cannot be tampered by malicious accesses.

4.3.1 Monitoring IP

The first operation performed by the architecture presented in Figure 10 is the monitoring of attack events thanks to a dedicated IP (labeled "Monitoring IP"in Figure 10). In order to detect attacks, three flags are extracted from firewalls (both *Local* and *Cryptographic*):

- *checkFlag* (cF): one of the controls (read/write access right or data format) failed in the *Checking Module*. This signal is similar to the *check_out* signal described in Section 3.
- *notFoundFlag* (nF): attempt to access an unknown address (i.e. an address not included in the *Correspondence Table*).
- *authenticationFlag* (aF): authentication verification failed (flag specific to a *Cryptographic Firewall* linked to the external memory controller).

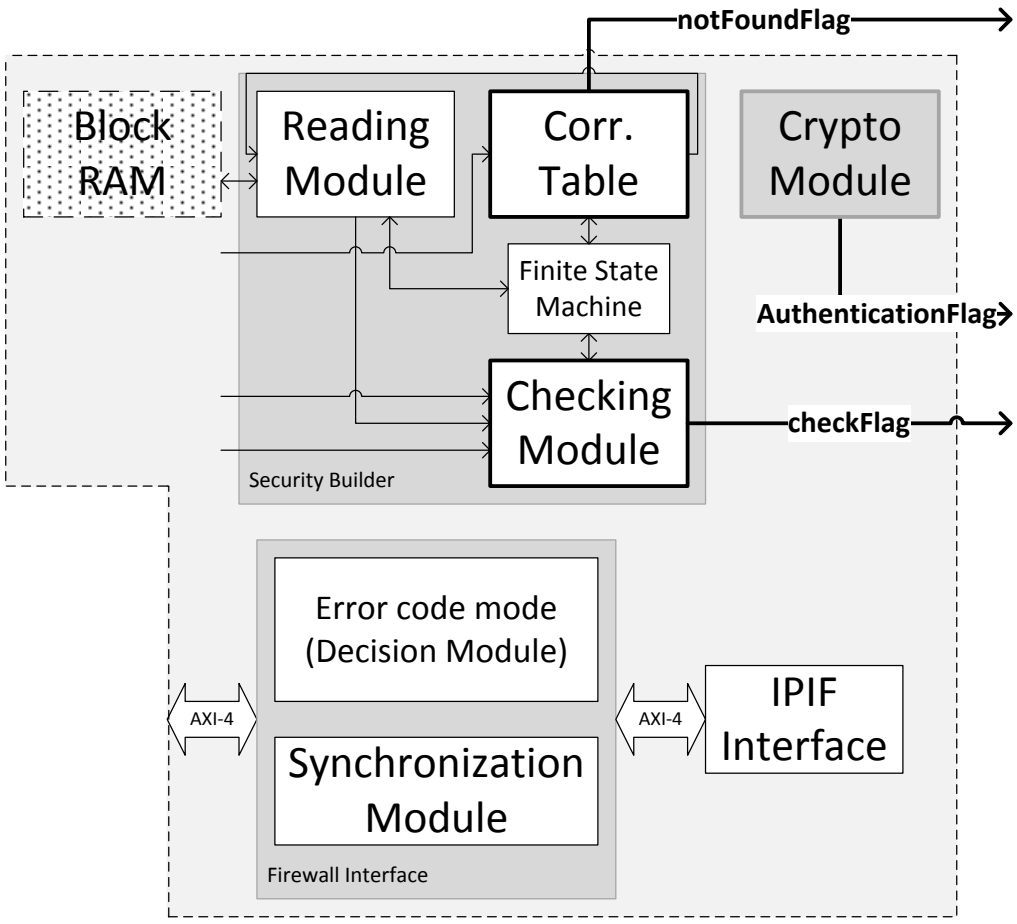

Figure 11: Generic structure of a firewall with flags extraction

Figure 11 shows a mixed structure of a *Local Firewall* and a *Cryptographic Firewall*. Each flag is stored on a significant bit in internal registers of the monitoring IP (each IP has a dedicated register).

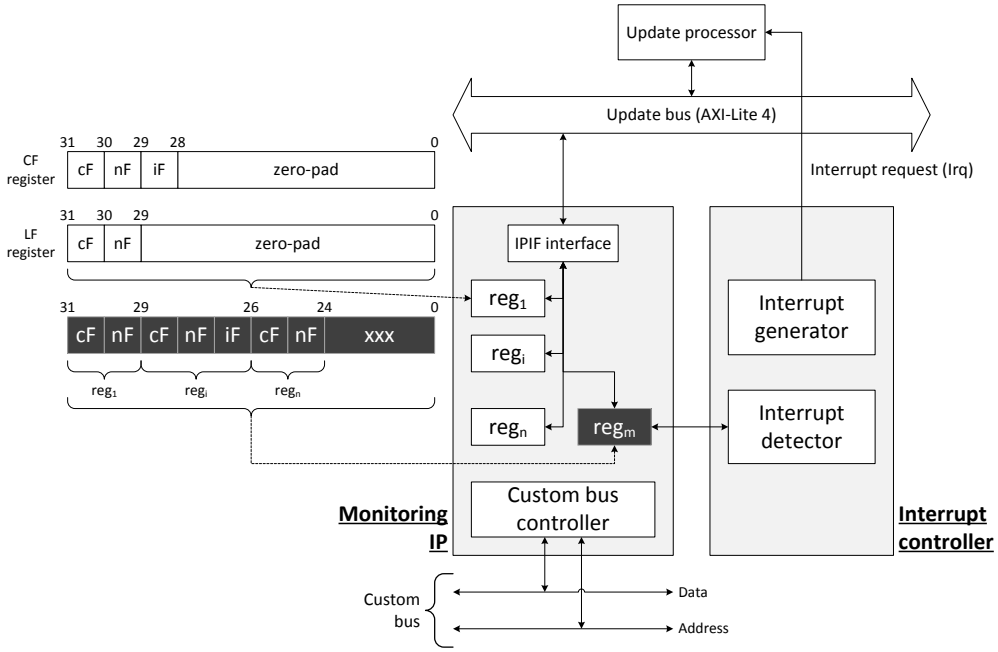

Figure 12: Architecture of the monitoring IP and interrupt routine

Figure 12 shows the detailed structure of the monitoring IP. Each firewall provides its flag information through a custom bus (address and data). Each reg_i is a 32-bit register with only 3 significant bits (corresponding to 3 flags). The main register reg_m concatenates all the significant bits. As reg_m register is also on 32 bits, the monitoring IP can manage up to 10 firewall registers. An interruption controller reads the main register reg_m and sends an interrupt request to the update processor as soon as one of the reg_m bits is equal to zero (i.e. when an error is detected by a firewall). This update processor (executing the interrupt routine) has several features:

- This processor is aware of the security context: security modes (see Section 3) are known for each firewall.
- When a firewall has to be updated, security policies are set to a less permissive configuration.

Beyond these monitoring tasks, this processor is also in charge of the security policies update (i.e. writing new values in Block RAMs linked to firewalls).

4.3.2 Security update protocol

Compared to the static protection described in Section 3, firewall structure is slightly modified. Both Block RAM data ports are used (one directly connected to the firewall, the other connected to the update processor). Therefore, it may happen that both data ports are used at the same time to access the same memory location. The architecture presented in Figure 13 has a mechanism, based on communication properties to avoid such cases.

MPSoCs considered in this work use the AXI protocol from ARM as the communication bus. Transactions in AXI protocol are based on a "handshake" protocol with couples of *valid*/*ready* signals for address and data channels (for both read and write [26]). Output information is only available when both signals (*valid* and *ready*) are in their high state (no information is transmitted in other cases). All the transactions using the AXI communication protocol are performed in two steps: a handshake on addresses is performed before a handshake on data signals.

Therefore, a mechanism is implemented in each *Firewall Interface* module to avoid concurrent accesses to BRAM memory location by both firewall and update processor. When an attack is detected in a firewall, output ready signals of the *Firewall Interface* are kept in their low state.

This step aims to "freeze" the communication bus and prevent malicious accesses during the update process. Then, the update process is executed in accordance with the simplified architecture shown in Figure 13.

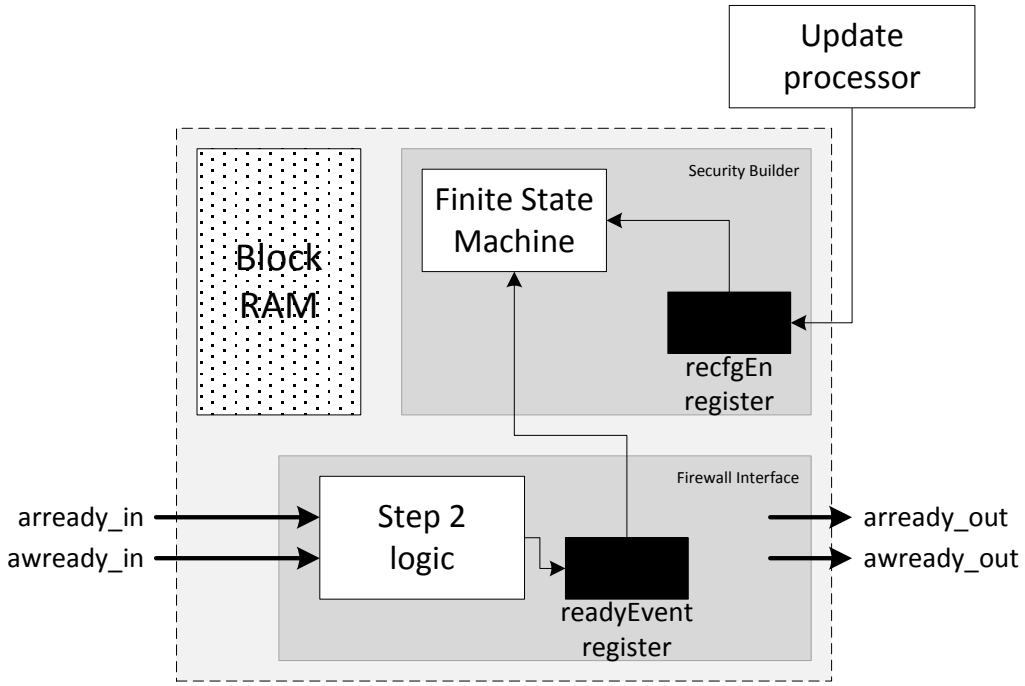

Figure 13: Partial architecture of a firewall with update process information

First, security update is enabled by the update processor writing a "1" value in the *recfgEn* register embedded in the *Security Builder*. Thus, the FSM associated to this process goes from a monitoring state to an update state. Then, while the processor retrieves the new security policy to be written into the memories, all transactions are blocked as previously described using the "handshake" feature. If a rising edge occurs on a ready signal (*arready_in* or *awready_in*), it is blocked until the update process is finished (i.e. $arready_out = 0$ and $awready_out = 0$ while $recfgEn = 1$) and a "1" value is written into the *readyEvent* register embedded in the *Firewall Interface* module. When the new security policy has been written, the update processor writes a "0" value in the *recfgEn* register to notify the end of the process.

Finally, according to the *readyEvent* register value, data blocked in the input port of the Firewall Interface is analyzed with the new security policy only if $readyEvent = 1$ (i.e. if a rising edge on a ready signal has been detected during update). The overall latency of the update process depends on the number of security policies to be modified. This process usually contains five steps:

- Extracting flags from firewalls to the monitoring IP and blocking of firewalls. This step is done in 1 clock cycle.
- Interruption routine execution. An interrupt request is sent in 2 cycles.
- Computation of the new security configuration (performed in software). It corresponds to the software latency on the update processor. For a basic implementation, the new security policy is retrieved in 148 clock cycles. Nevertheless, other algorithms may be considered. In this case, response time would be impacted.
- Writing the new security policy. This step depends of the number of security policies to be written. Results are given in Section 5.
- Reactivation of the main application after update completeness. This is done in 1 clock cycle.

5 Implementation results

5.1 Experimental setup

In order to validate this contribution, implementations and simulations are done on a case study representative of a real embedded system on a Xilinx ML605 development board (including a Virtex-6 FPGA). We have developed a system performing image processing operations. Its architecture is shown in Figure 14, it contains 2 MicroBlaze softcore processors and 2 IPs:

- An image processing IP. This IP contains several programmable registers and performs a threshold function on a picture.
- It is assumed that temporary pictures and other information (code, user profiles and so on) are stored in a shared Block RAM memory.

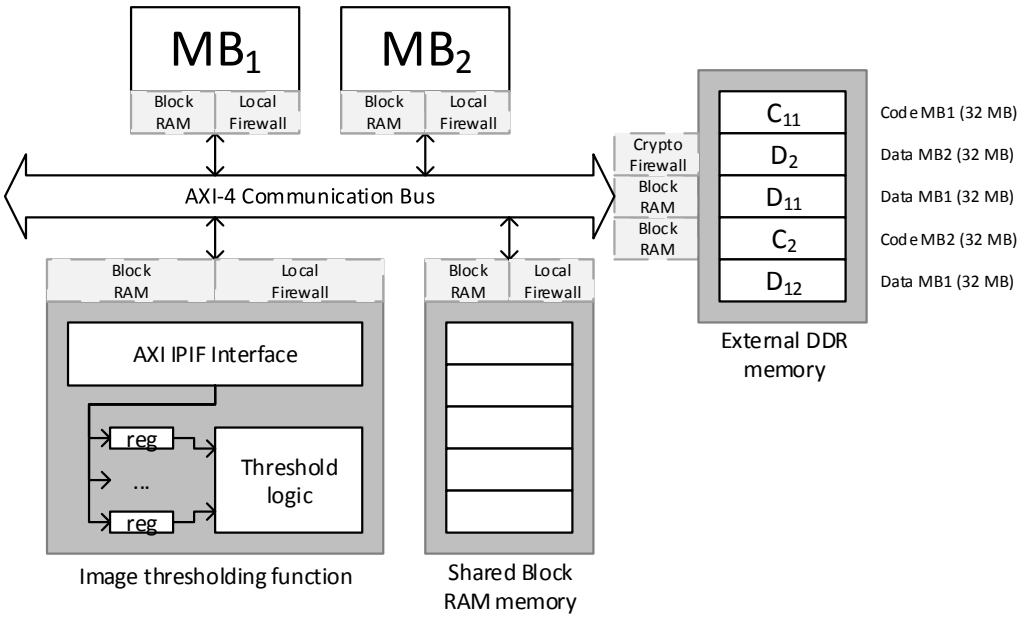

Figure 14: Case study architecture

It is also assumed that the external memory is shared between the 2 processors and split into memory sections of 32 MB each:

- Processor MB1: 1 code section (C11) and 2 data sections (D11 and D12).
- Processor MB2: 1 code section (C21) and 1 data section (D21).

Regarding cryptographic options, memory sections can be protected in "confidentiality and integrity" (C11 and D11), "integrity only" (D12) or even in plaintext (C21 and D21). All transactions are based on a 32-bit data format. Finally access rights are defined in Table 3.

Table 3: Access rights for the case study

	Shared memory	**Image processing**	**External memory**	
			C21, D21	C11, D11, D12
MB1	Read only	Read/write	No access	Read/write
MB2	Read/write	Write only	Read/write	No access

These security policies are built to illustrate all the cases (in terms of policies and cryptographic options) on a multiprocessor architecture implemented on a FPGA chip. Then, the case study has some requirements on the software layer (for both applications and operating systems):

- On the one hand, there are some requirements on the operating system. The *standalone* OS provided by Xilinx development tools is used: it does not take into account multitask aspects but allows making measurements on communications between elements of the system. Consequently, a POSIX[6] compliant operating system is used to run custom applications *picProc*, *picDRM* and *picDec*. Xilkernel (from Xilinx) is chosen but *mu*COS[7] or *mu*CLinux[8] could have been considered as well.

- On the other hand, custom applications are developed to illustrate all the situations (access to memories, communications between processors...).

The main application *picProc* is based on a sample JPEG image. All operations are saved in a log file. *picProc* is executed through the following steps:

- First, MB_1 reads the image from the external memory ($m_encryptPic$) and processes a software deciphering.

- Then, the plaintext image $m_decryptPic$ is stored in the shared BRAM memory.

- The next step is the transfer of the plaintext image to the threshold IP. This is performed again by processor MB_1.

- Finally, processor MB_2 writes the threshold image ($m_processedPic$) in plaintext into the external memory.

The picProc application process is summarized in Figure 15.
Beyond that, other applications are used to evaluate the case study:

- *picDRM* is an application ran by a processor checking DRM rights of the target image while the other processor does read/write operations.

- *picDec* is an application performing a software deciphering of an image by a unique processor. It serves as a reference for latency results as it does not imply a *Cryptographic Firewall*.

[6]http://standards.ieee.org/develop/wg/POSIX.html
[7]http://micrium.com/rtos/ucosii/overview/
[8]http://www.uclinux.org/

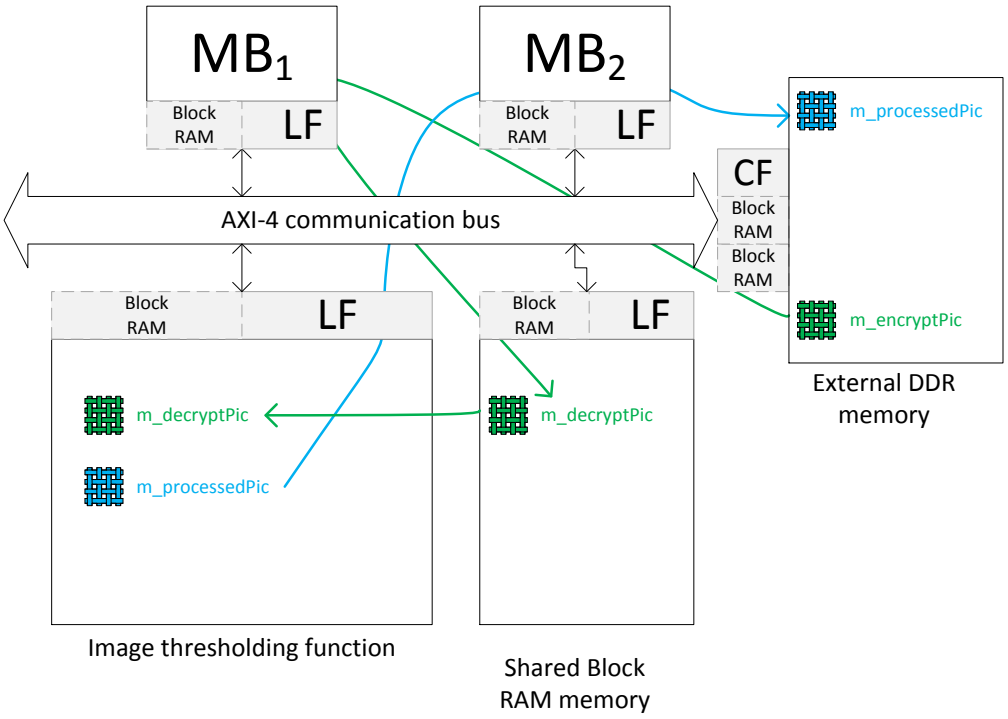

Figure 15: Illustration of picProc on the case study

5.2 Security analysis

5.2.1 Context

The main goal of this section is to provide a security analysis for this work compared to other approaches [17, 21]. This analysis is based on the case study described in the previous section where two attack scenarios will be discussed to verify how implementations behave against the threat model defined in Section 2. Implementations of [17] and [21] on the previously defined case study is described in Figure 16.

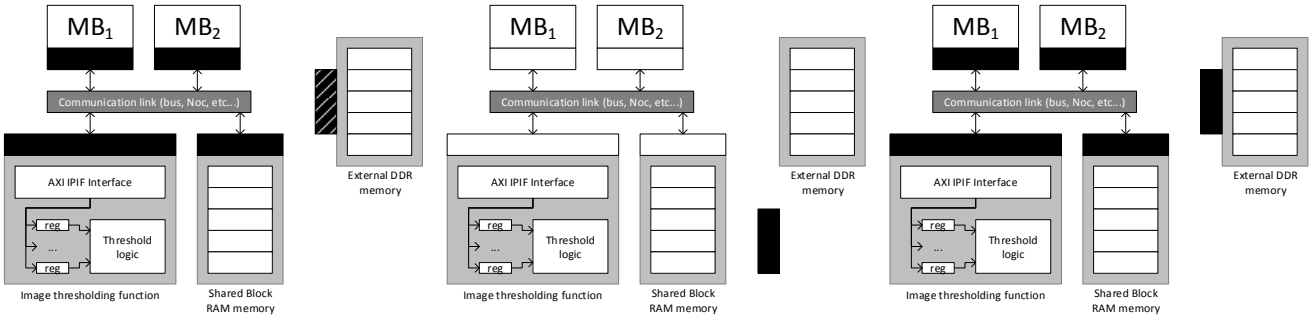

(a) This work: distributed cryptography-enhanced memory inter-face

(b) [21] approach: security checking inter-centralized in a specific module

(c) [17] approach: distributed with a generic memory interface

Figure 16: Distributed and centralized behaviors

In Figure 16, it is assumed that the communication network is the one used in each reference. This is due to the fact that in [17, 21] security mechanisms rely partially on the communication architecture. In the original case study, there are memory pages with several cryptographic features.

This security analysis explores three options for our approach:

- Without cryptographic features (to be fair with existing works that do not provide such services). Therefore, all the data is in plaintext. It would be similar to the case where only *Local Firewalls* are implemented.
- With integrity only.
- With both confidentiality and integrity (the full AES-GCM core is enabled).

Then, the two scenarios considered in this analysis are as follows:

- Scenario S1. In a OCR (*Optical Character Recognition*) context, an attacker changes a plaintext data packet with another. For example, a packet containing an image and the address of the next packet is replaced with a noised version and an unknown address. In this case, the threshold function computes an unexpected result.
- Scenario S2. Actually, the threshold value is hardcoded in the IP. It is assumed that the attacker was able to either build another IP with another threshold value or reconfigure the system with his malicious IP. As the threshold is different, the result will lead to different results: especially in a OCR context, this attack may dramatically change the application performance.

Following subsections described how implementations of existing works behave according to these scenarios.

5.2.2 Scenario S1

With integrity only. In this case, the modified data is detected by the *Cryptographic Firewall*.

With confidentiality and integrity. Beyond the fact that data is ciphered, the attack is detected by the *Cryptographic Firewall* with integrity checking.

Without cryptographic services. For this work, it is considered that a simple *Local Firewall* is attached to the external memory controller for our approach. For both [17] and this work, the malicious packet will be correctly read by processor MB1. Packet will be blocked only at the communication interface of MB1. For [21], even if the process is similar, the attack is detected later at the security manager layer which is decentralized in a specific IP. For the second stage of this attack (transmission to a corrupted address), [17] and this work still detect it at MB1 bus interface while [21] detects it even later within its security manager.

5.2.3 Scenario S2

This scenario is different as a malicious IP is implemented in the system. Therefore, cryptographic features that can be applied to this work (AES-GCM core in the *Cryptographic Firewall*) are not relevant. [17] and [21] behaves as the previous attack. This work is more relevant as the malicious IP does not provides a *Local Firewall*, monitoring process will be stopped earlier (as soon as data comes in the threshold IP interface): in fact, the update processor (described in Section 4) is also responsible for monitoring firewalls states. As a consequence, the update processor is aware that the threshold function is performed by an untrusted entity.

5.2.4 XBox 360 use case

Security in game consoles has been a hot topic since the 90s with the introduction of modchips for the first generation of Playstation [33]. The XBox 360 is a game console developed by Microsoft in 2005. It was one of the first consoles to take into account existing vulnerabilities aiming to take the control of the platform. XBox 360 was designed to include several security functions. The default memory layout of such a console is presented in Figure 17.

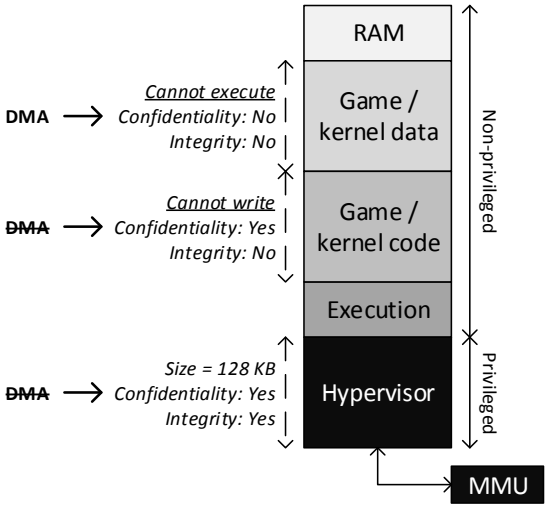

Figure 17: Xbox 360 memory layout

RAM memory is divided in four sections: game/kernel data, game/kernel code, execution section (reserved for reads and writes during runtime) and hypervisor code. Furthermore, the game/kernel data section (which is in plaintext) must be accessed through DMA for specific purposes (basically GPU-related operations) while other memory sections are secured (at least, in terms of confidentiality). In this default configuration, data section is a clear security breaches which can lead to critical exploits such as the "King Kong" attack: this is a well-known method where the attacker uses an event where the GPU access data section in order to get hypervisor-like privileges. In case the system is compromised, the system could accessed unexpected sections. If *Local Firewalls* are implemented in such a system, even if the memory is compromised by one of its memory sections, the approach described in this work could detect the attack in two times:

- As security rules are written in secured on-chip memories, it is considered that accesses to privileged sections can be blocked.
- Even if an hypervisor-like user perfoms unwanted operations on other components of the console architecture, abnormal behaviors are detected through *Local Firewalls*.

Given the fact that console developers are able to implement their architectures on a FPGA chip, the solution described in this work is an efficient solution to block malicious exploits at both hardware and software levels.

5.3 Experimental results for static and dynamic security

5.3.1 Area results

This study is about firewall areas compared to a reference (a MicroBlaze softcore processor running at 100MHz without caches). Firewalls structure was described in Section 3. The *Correspondence Table* is the only module with a relation between area and security (3 registers are needed to define each security policy).

It is assumed that the following results consider the worst case (i.e. the largest *Correspondence Table* with 10 security policies). Table 4 shows area results of both firewalls (*Local* and *Cryptographic*) compared to the reference MicroBlaze processor.

Table 4: Area results of firewall implementation

		Slices	**Slice regs**	**LUTs**	**# of Block RAMs**
Local Firewall	**Firewall Interface**	76	120	68	0
	Security Builder	23	3	55	1
	Total	99	123	293	1
Crypto Firewall	**Firewall Interface**	76	120	153	0
		23	3	55	1
	Crypto Module	1,166	2,038	2,396	14
		89.42%	94.31%	89.10%	93.33%
	Total	1,304	2,161	2,689	15
MicroBlaze		1,179	1,298	1,829	10

One *Local Firewall* has an acceptable area cost compared to one *Cryptographic Firewall* (around 9%) and to one Microblaze (in this case, nearly 11%). This is essentially due to the AES-GCM ciphering core: its area corresponds to 90% of the overall area of one *Cryptographic Firewall*. Therefore, we can also estimate the area of a complete system embedding x *Local Firewalls* and y *Cryptographic Firewalls*. It is assumed that synthesis tools are linear: the area of a structure with 2 identical modules is twice the area of a single module. Areas in terms of slices, registers and LUTs are given in following equations:

$$numSlices = 138 \times x + 1,304 \times y \tag{3}$$

$$numRegs = 123 \times x + 2,161 \times y \tag{4}$$

$$numLuts = 293 \times x + 2,689 \times y \tag{5}$$

Designer can use these equations to estimate the cost of the hardware firewalls for his/her system. Table 5 shows area overheads due to the update mechanisms presented in Section 4.

Table 5: Synthesis results for the update mechanisms on a Local Firewall

		Slices	**Slice regs**	**LUTs**	**# of BRAMs**
Static solution (Local Firewall)		138	123	293	1
Improvements for update solutions	**Real time**	6	0	5	0
	Quarantine	17	0	18	0
	Total overhead	5	13	15	0
	Total overhead	+20.29%	+10.57%	+12.97%	+0.00%

Table 5 only shows modifications on a firewall (a complete solution would include a processor, a communication bus...). The "improvements for update"contribution corresponds to flip-flops used to block the data during the update process (see Section 4). Otherwise, extensions on a *Local Firewall* leads to a 20% area overhead (it would be relatively lower on a *Cryptographic Firewall*). Table 6 shows area overheads described in Section 5.1.

Table 6: Area results of several configurations

	Slices	**Regs**	**LUTs**	**# of BRAMs**
Solution x0 without firewalls	5,446	7,195	8,354	32
Solution x1a with firewalls and without updates	7,302 +34.08%	9,848 +36.87%	12,215 +46.22%	51 +37.25%
Solution x1b with firewalls and update	7,442 +1.92%	9,913 +0.66%	12,405 +1.55%	51 +0.00%
Solution x2a with firewalls and without update	+23.38%	+22.71%	+32%	+44.25%

Four configurations are taken into account:

- $x0$: this is the original case study without firewalls as shown in Figure 14.
- $x1a$: $x0$ architecture with static firewalls as described in Section 3.
- $x1b$: $x1a$ architecture with update improvements (Section 4).
- $x2$a: twice $x1a$ architecture with static firewalls as described in Section 3: 4 processors, 4 IPs but still one external memory. This configuration is used to illustrate the scalability of the firewalls approach described is this work.

Case study with static firewalls implies a significant area overhead (34% in terms of slices): this is mainly due to the AES-GCM core embedded in the *Cryptographic Firewall*. Logic added for update purposes implies an overhead less than 2% compared to the static version. Regarding the scalability of this approach, the area overhead is up to 24% in terms of slices for a case study twice larger (configuration $x2a$). It must be noted that the area overhead would proportionally decrease as the architecture grows up (in terms of percentages). Finally, it is possible to make a quantitative and qualitative comparison with [16, 21]. Table 7 presents some results.

Table 7: Comparison with existing efforts

		[21]	**[17]**	**[13]**	**[9]**	**This work**
Standalone monitoring block	**Processor ratio (%)**	SEI	DPU	Full	SMM + Ctrl	LF
		6.20%	25%	100%	73%	11.30%
	Reference (LUTs)	N/A	N/A	1829	947	293
Case study overhead	**Reference**	N/A	N/A	Nude	+1 SMM/Ctrl	+4 LF
				+0%	+11.34%	+14.03%
	Custom	+1 LF	+4 LF	Nude	+1 LF	+4 LF
		+11.30%	+14.03%	+0%	+11.30%	+14.03%
Update features		No	Yes	No	No	Yes

In terms of area, this work is less efficient than the SECA solution [21]: we propose an approach where controls are performed in each firewall while SECA has a centralized security checker. Otherwise, our approach is better than [16] solution: for security update, no additional memory is needed as a communication property is used to block data ("handshake" feature of AXI protocol) while Fiorin et al. approach stores all incoming data in a buffer until update completeness.

5.3.2 Latency results

In this section we propose to analyze the latency of realistic scenarios based on applications described in Section 5.1 (*picProc*, *picDRM* and *picDec*). Simulations were done with the following scenarios:

- S0: latency of a single *Local Firewall*.
- S1: communication between two elements within the architecture (use of two *Local Firewalls*).
- S2: communication between a processor and an external memory section protected with confidentiality and integrity (implies one *Local Firewall* and a *Cryptographic Firewall*).
- S3: communication between a processor and an integrity only section in the external memory.
- S4: communication between a processor and a plaintext section in the external memory.

Latency of a *Local Firewall* (scenario S0) serves as a reference, other scenarios imply two firewalls (two *Local* for S1 while S2, S3 and S4 use one *Local* and one *Cryptographic Firewall*). Results are given in Figure 18. In terms of latency, the most critical scenario is S2 as it implies cryptographic

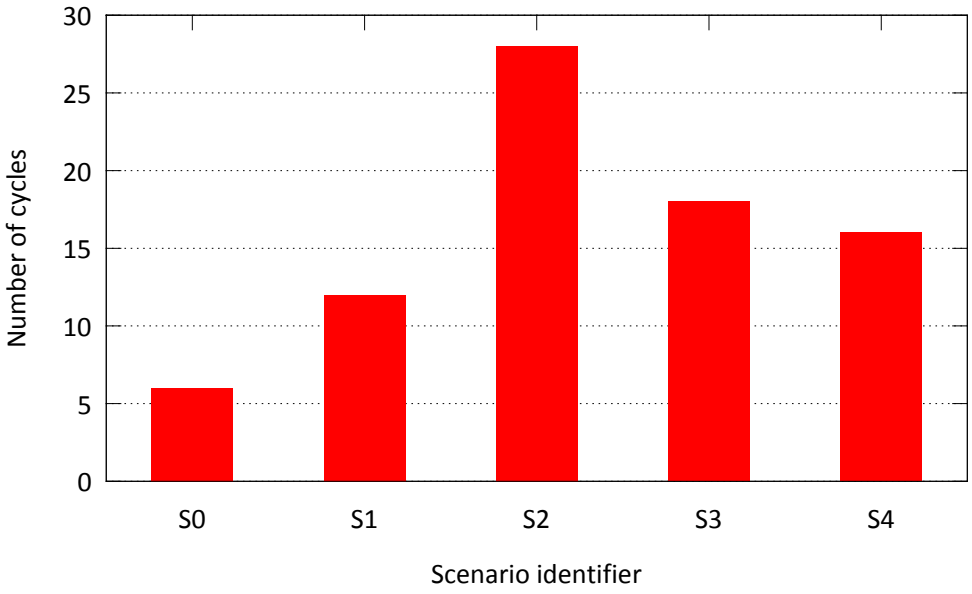

Figure 18: Latency results of scenarios

operations (access to the external memory). S2 lasts for 28 clock cycles:

- **6 cycles** for the security policy checking as in a *Local Firewall*, identically to scenario S0.
- According to AES-GCM, protection with confidentiality/integrity of N data blocks is performed in $10 + (10 + 2) \times N$ cycles (therefore, **22 cycles** for a single data block).

Case study with 8KB caches (both instruction and data) is considered in this study. It is assumed that the overall execution time is divided in two parts:

- Accesses to the external memory (cache miss or read/write operations). It implies 1 *Cryptographic Firewall* and 1 *Local Firewall*.
- Internal accesses (between a processor and an IP or between two processors). This option only implies *Local Firewalls*.

Table 8: Latency overheads for custom applications

	Execution time (ms)	**# of accesses to external memory**	**Latency overhead [17] and this work approach**
picProc	3,623	34,063,298	17.76%
picDrm	1,084	9,642,055	9.43%
picDev	382	4,736,966	4.18%

Timers and processor specific registers are used to extract information and compare latency overheads. The simplest application (*picDec*) has an overhead lower than 5% (it does not require cryptographic features). *picProc* is the application with the highest penalty (nearly 18%) as it implies a *Cryptographic Firewall* performing AES-GCM operations. In order to compare this work with existing solutions, firewalls approach is transposed in a centralized approach as it is proposed in the SECA model [21].

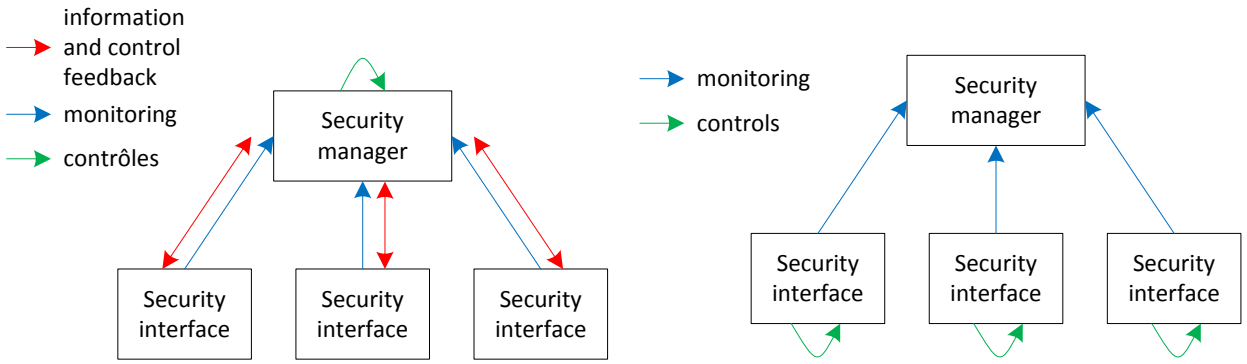

(a) Distributed (b) Centralized

Figure 19: Distributed and centralized behaviors

In a distributed approach (as firewalls, see Figure 19a), security controls are performed locally and monitoring information is sent to the security manager when an attack is detected (this is done by the monitoring IP and the update processor). In a centralized approach ([21], see Figure 19b), each security interface sends information to the security manager doing all the security controls. These connections are considered as trusted areas, but centralization implies latency overhead for communications between each interface and the manager. This work considers that each transaction (a roundtrip between an interface and a manager) through an AXI-Lite bus requires 4 additional clock cycles [26, 34]. Therefore, for the previously defined *picDec* application, a centralized implementation gives a 6.18% latency overhead while the distributed firewall approach has a 4.18% overhead (a 33% gain compared to SECA approach).

5.3.3 Memory occupancy

According to firewall description detailed in Section 3, internal memories can be used to store:

- Security policies: parameters to be verified by firewalls.
- Timestamps: used by the AES-GCM core to protect against replay attacks. It is assumed that timestamps are generated with internal counters.
- Tags: when an external memory section has to be protected in terms of integrity, a tag is needed.

This work only focuses on tags produced by the AES-GCM core. For each 128-bit data, a tag of equal length is stored. Therefore, we can only protect 1.87MB of data in terms of integrity in a Xilinx ML605 board (XC6VLX240T1156-1 FPGA).

- At first sight, this can be a limiting factor. A more efficient approach would be the use of Bloom filters [35, 36]: under specific conditions, this technique can decrease the memory footprint by 96%.
- However, firewalls offer flexibility in terms of system security. In accordance with user requirements, data may not always be protected in terms of integrity (firewalls are able to detect other threats).

6 Conclusion and future work

Several related works show an interest in communication and memory protection. Nevertheless, no approach has addressed the need for cryptographic functions aiming at protecting memories with monitoring and controls function providing communication protection. The distributed firewall approach allows embedded system developers to protect both communications and memories in a multiprocessor architecture.

Firewalls implementation shows that security can be provided only with hardware firewalls. The additional software is only for update purposes (attack detection, new configurations computation...). This work proposes a low-latency solution where latency overheads can reach 17% (according to cryptographic features set in firewalls).

One perspective regarding our firewalls approach is a software protection: firewalls would analyze data according to the current software task identifier, which allows refining the security level of the overall solution. Finally, even if hardware security is a recent domain, it is possible to implement efficient functions; we believe that technology improvements (for instance, multi-layer FPGAs or hybrid platforms embedding FPGAs and processors) will allow going further in protection of multiprocessor architectures.

References

[1] P. Kocher, R. Lee, G. McGraw, and A. Raghunathan, "Security as a new dimension in embedded system design," in *Proceedings of the 41st Annual Design Automation Conference*, DAC '04, (New York, NY, USA), pp. 753–760, ACM, 2004. Moderator-Ravi, Srivaths.

[2] S. Ravi, A. Raghunathan, P. Kocher, and S. Hattangady, "Security in embedded systems: Design challenges," *ACM Trans. Embed. Comput. Syst.*, vol. 3, pp. 461–491, Aug. 2004.

[3] D. Lie, C. A. Thekkath, and M. Horowitz, "Implementing an untrusted operating system on trusted hardware," *SIGOPS Oper. Syst. Rev.*, vol. 37, pp. 178–192, Oct. 2003.

[4] G. Suh, C. O'Donnell, and S. Devadas, "Aegis: A single-chip secure processor," *Design Test of Computers, IEEE*, vol. 24, pp. 570–580, Nov 2007.

[5] L. Bossuet, M. Grand, L. Gaspar, V. Fischer, and G. Gogniat, "Architectures of flexible symmetric key crypto engines—a survey: From hardware coprocessor to multi-crypto-processor system on chip," *ACM Comput. Surv.*, vol. 45, pp. 41:1–41:32, Aug. 2013.

[6] A. J. Mahar, P. M. Athanas, S. D. Craven, J. N. Edmison, and J. Graf, "Design and characterization of a hardware encryption management unit for secure computing platforms," in *39th*

Hawaii International International Conference on Systems Science (HICSS-39 2006), CD-ROM / Abstracts Proceedings, 4-7 January 2006, Kauai, HI, USA, 2006.

[7] R. Elbaz, L. Torres, G. Sassatelli, P. Guillemin, M. Bardouillet, and A. Martinez, "A parallelized way to provide data encryption and integrity checking on a processor-memory bus," in *Proceedings of the 43rd Annual Design Automation Conference*, DAC '06, (New York, NY, USA), pp. 506–509, ACM, 2006.

[8] R. Vaslin, G. Gogniat, J.-P. Diguet, R. Tessier, D. Unnikrishnan, and K. Gaj, "Memory security management for reconfigurable embedded systems," in *ICECE Technology, 2008. FPT 2008. International Conference on*, pp. 153–160, Dec 2008.

[9] J. Crenne, R. Vaslin, G. Gogniat, J.-P. Diguet, R. Tessier, and D. Unnikrishnan, "Configurable memory security in embedded systems," *ACM Trans. Embed. Comput. Syst.*, vol. 12, pp. 71:1–71:23, Apr. 2013.

[10] J. Crenne, P. Cotret, G. Gogniat, R. Tessier, and J.-P. Diguet, "Efficient key-dependent message authentication in reconfigurable hardware," in *Field-Programmable Technology (FPT), 2011 International Conference on*, pp. 1–6, Dec 2011.

[11] M. Haifeng, C. Jie, and G. Zhenguo, "Memory integrity protection method based on asymmetric hash tree," in *Distributed Computing and Applications to Business, Engineering and Science (DCABES), 2014 13th International Symposium on*, pp. 263–267, Nov 2014.

[12] T. Wiersema, S. Drzevitzky, and M. Platzner, "Memory security in reconfigurable computers: Combining formal verification with monitoring," in *Field-Programmable Technology (FPT), 2014 International Conference on*, pp. 167–174, Dec 2014.

[13] Xilinx, "Microblaze processor reference guide - embedded development kit edk 14.7," tech. rep., Xilinx, 2012.

[14] J.-P. Diguet, S. Evain, R. Vaslin, G. Gogniat, and E. Juin, "Noc-centric security of reconfigurable soc," in *Networks-on-Chip, 2007. NOCS 2007. First International Symposium on*, pp. 223–232, May 2007.

[15] L. Fiorin, C. Silvano, and M. Sami, "Security aspects in networks-on-chips: Overview and proposals for secure implementations," in *Digital System Design Architectures, Methods and Tools, 2007. DSD 2007. 10th Euromicro Conference on*, pp. 539–542, Aug 2007.

[16] L. Fiorin, S. Lukovic, and G. Palermo, "Implementation of a reconfigurable data protection module for noc-based mpsocs," in *Parallel and Distributed Processing, 2008. IPDPS 2008. IEEE International Symposium on*, pp. 1–8, April 2008.

[17] L. Fiorin, G. Palermo, and C. Silvano, "A security monitoring service for nocs," in *Proceedings of the 6th IEEE/ACM/IFIP International Conference on Hardware/Software Codesign and System Synthesis*, CODES+ISSS '08, (New York, NY, USA), pp. 197–202, ACM, 2008.

[18] J. Sepulveda, R. Pires, G. Gogniat, W. J. Chau, and M. Strum, "Qoss hierarchical noc-based architecture for mpsoc dynamic protection," *Int. J. Reconfig. Comput.*, vol. 2012, pp. 3:3–3:3, Jan. 2012.

[19] J. Sepulveda, G. Gogniat, C. Pedraza, R. Pires, W. J. Chau, and M. Strum, "Hierarchical noc-based security for mp-soc dynamic protection," in *Circuits and Systems (LASCAS), 2012 IEEE Third Latin American Symposium on*, pp. 1–4, Feb 2012.

[20] M. LeMay and C. A. Gunter, *Logic, Rewriting, and Concurrency: Essays Dedicated to José Meseguer on the Occasion of His 65th Birthday*, ch. Network-on-Chip Firewall: Countering Defective and Malicious System-on-Chip Hardware, pp. 404–426. Cham: Springer International Publishing, 2015.

[21] J. Coburn, S. Ravi, A. Raghunathan, and S. Chakradhar, "Seca: Security-enhanced communication architecture," in *Proceedings of the 2005 International Conference on Compilers, Architectures and Synthesis for Embedded Systems*, CASES '05, (New York, NY, USA), pp. 78–89, ACM, 2005.

[22] T. Huffmire, B. Brotherton, N. Callegari, J. Valamehr, J. White, R. Kastner, and T. Sherwood, "Designing secure systems on reconfigurable hardware," *ACM Trans. Des. Autom. Electron. Syst.*, vol. 13, pp. 44:1–44:24, July 2008.

[23] G. Plouviez, *Étude, spécification, vérification formelle de mécanismes de virtualisation sécurisés pour architecture many-cores*. PhD thesis, Université Pierre et Marie Curie (Paris, France), 2011.

[24] J. PÃ¶rquet, *Architecture de sécurité dynamique pour systèmes multiprocesseurs intégrés sur puce*. PhD thesis, Université Pierre et Marie Curie (Paris, France), 2010.

[25] Z. Yan-ling, P. Wei, and Z. Xin-guo, "Design and implementation of secure embedded systems based on trustzone," in *Embedded Software and Systems, 2008. ICESS '08. International Conference on*, pp. 136–141, July 2008.

[26] ARM, "Amba axi and ace protocol specification axi3, axi4, and axi4-lite, ace and ace-lite," tech. rep., ARM, 2012.

[27] Xilinx, "Zynq-700 epp overview," tech. rep., Xilinx, 2012.

[28] ARMadeus, "Embedded linux board - designed by armadeus systems," tech. rep., ARMadeus, 2014.

[29] K. Jarvinen, M. Tommiska, and J. Skytta, "Hardware implementation analysis of the md5 hash algorithm," in *System Sciences, 2005. HICSS '05. Proceedings of the 38th Annual Hawaii International Conference on*, pp. 298a–298a, Jan 2005.

[30] Z. Shi, C. Ma, J. Cote, and B. Wang, "Hardware implementation of hash functions," in *Introduction to Hardware Security and Trust*, Springer, 2012.

[31] R. Chaves, G. Kuzmanov, L. Sousa, and S. Vassiliadis, "Improving sha-2 hardware implementations," in *Cryptographic Hardware and Embedded Systems - CHES 2006* (L. Goubin and M. Matsui, eds.), vol. 4249 of *Lecture Notes in Computer Science*, pp. 298–310, Springer Berlin Heidelberg, 2006.

[32] D. McGrew and J. Viega, "The security and performance of the galois/counter mode (gcm) of operation," in *Progress in Cryptology - INDOCRYPT 2004* (A. Canteaut and K. Viswanathan, eds.), vol. 3348 of *Lecture Notes in Computer Science*, pp. 343–355, Springer Berlin Heidelberg, 2005.

[33] E. DeBusschere and M. McCambridge, "Modern game console exploitation," tech. rep., Department of Computer Science, University of Arizona, 2012.

[34] N.-C. Chang, Y.-Z. Liao, and T.-S. Chang, "Analysis of shared-link axi," *Computers Digital Techniques, IET*, vol. 3, pp. 373–383, July 2009.

[35] B. H. Bloom, "Space/time trade-offs in hash coding with allowable errors," *Commun. ACM*, vol. 13, pp. 422–426, July 1970.

[36] J. Crenne, *Sécurité haut-débit pour les systèmes embarqués à base de FPGAs*. PhD thesis, Université de Bretagne-Sud (Lorient, France), 2011.